

Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals - group of 52 »

[All articles](#) [Recent articles](#)

J Gray, S Chaudhuri, A Bosworth, A Layman, D ... - Data Mining and Knowledge Discovery, 1997 - Springer
... and roll-up operators along with system and **user-defined** aggregates are ... Beyond the
five standard **aggregate functions** defined so far, many **SQL** systems add sta ...
Cited by 1017 - [Web Search](#) - [BL Direct](#)

Providing database as a service - group of 9 »

H Hacigumus, B Iyer, S Mehrotra - Data Engineering, 2002. Proceedings. 18th International ..., 2002 - ieeexplore.ieee.org
... the no- tation for a simple select query with **user defined functions, aggregate functions**, GROUP BY ... as μ , where is one of the aggregation **functions** of **SQL**. ...
Cited by 55 - [Web Search](#) - [BL Direct](#)

On parallel processing of aggregate and scalar functions in object-relational DBMS - group of 6 »

M Jaedicke, B Mitschang - ACM SIGMOD Record, 1998 - portal.acm.org
... 3.1 Built-in and **User-Defined Functions** Every RDBMS comes ... **functions** can be either
scalar **functions** or **aggregate functions**. A scalar **function** can be used in **SQL** ...
Cited by 25 - [Web Search](#) - [BL Direct](#)

An efficient pictorial database system for PSQL - group of 6 »

N Roussopoulos, C Faloutsos, T Sellis - IEEE Transactions on Software Engineering, 1988 - doi.ieeecomputersociety.org
... extension of **SQL** [2] and supports **user-defined** abstract data ... defined on a domain
may be simple or **aggregate**. A simple **function** receives elements from a set of ...
Cited by 132 - [Web Search](#) - [Library Search](#)

User defined aggregates in object-relational systems - group of 8 »

H Wang, C Zaniolo - Data Engineering, 2000. Proceedings. 16th International ..., 2000 - ieeexplore.ieee.org
... and execute using the (**SQL**-compliant) data ... can be built using the **aggregate functions**
found in ... mining algorithms require powerful **User-Defined Aggregates** (UDAs ...
Cited by 16 - [Web Search](#) - [BL Direct](#)

[PS] Independent, Open Enterprise Data Integration - group of 3 »

JM Hellerstein, M Stonebraker, R Caccia - IEEE Data Engineering Bulletin, 1999 - sites.computer.org
... marts are nothing more or less than **SQL** database systems ... standard relational language,
SQL99, extended with **user-defined** scalar and **aggregate functions**. ...
Cited by 45 - [View as HTML](#) - [Web Search](#)

Integrating Association Rule Mining with Relational Database Systems: Alternatives and Implications - group of 22 »

S Sarawagi, S Thomas, R Agrawal - Data Mining and Knowledge Discovery, 2000 - Springer
... a stored procedure; caching the data to a file system on-the-fly and mining;
tight-coupling using primarily **user-defined functions**; and **SQL** implementations for ...
Cited by 240 - [Web Search](#) - [Library Search](#) - [BL Direct](#)

A uniform model for temporal object-oriented databases - group of 4 »

GTJ Wu, U Dayal - Data Engineering, 1992. Proceedings. Eighth International ..., 1992 - ieeexplore.ieee.org
... and their **functions**. Complex (**aggregate**) types, such as sets, tuples, and multisets,
can be recursively built up from the primitive and **user-defined** types ...
Cited by 52 - [Web Search](#)

[BOOK] Deriving Production Rules for Constraint Maintenance - group of 7 »

S Ceri, J Widom - 1990 - stoev.org

... The core of the language is a variation on the usual **SQL** syntax for predicates ... I

!=|<|>|<=|>= sum 1 min 1 max 1 avg 1 count **user-defined aggregate function** +||*| ...

Cited by 246 - View as HTML - Web Search - Library Search

[\[PS\] Logic-Based User-Defined Aggregates for the Next Generation of Database Systems - group of 2 »](#)

C Zaniolo, H Wang - The Logic Programming Paradigm: Current Trends and Future ..., 1999 - magna.cs.ucla.edu

... system (languages with data-types compatible with **SQL**, such as Oracle PL/**SQL** and

JDBC ... In LDL ++, an **aggregate function** f on a set S, can be introduced via the ...

Cited by 18 - View as HTML - Web Search

Google ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

[Google Home](#) - [About Google](#) - [About Google Scholar](#)

©2006 Google

Terms used **user defined aggregate function sql**

Found **21,233** of **184,245**

Sort results by

 [Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

 [Search Tips](#)
[Try this search in The ACM Guide](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐


1 [On parallel processing of aggregate and scalar functions in object-relational DBMS](#)



Michael Jaedicke, Bernhard Mitschang

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.43 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Nowadays parallel object-relational DBMS are envisioned as the next great wave, but there is still a lack of efficient implementation concepts for some parts of the proposed functionality. Thus one of the current goals for parallel object-relational DBMS is to move towards higher performance. In this paper we develop a framework that allows to process user-defined functions with data parallelism. We will describe the class of partitionable functions that can be processed parallelly. We will ...

Keywords: aggregates, object-relational database systems, parallel query processing, user-defined functions

2 [Query processing for relational data: User-defined aggregate functions: bridging theory and practice](#)



Sara Cohen

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available:  pdf(267.75 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The ability to create user-defined aggregate functions (UDAs) is rapidly becoming a standard feature in relational database systems. Therefore, problems such as query optimization, query rewriting and view maintenance must take into account queries (or views) with UDAs. There is a wealth of research on these problems for queries with general aggregate functions. Unfortunately, there is a mismatch between the manner in which UDAs are created, and the information that the database system requires ...

Keywords: aggregate queries, query optimization, query rewriting, view maintenance

3 [Data Access and Knowledge Management: Advanced grouping and aggregation for data integration](#)




Eike Schallehn, Kai-Uwe Sattler, Gunter Saake

October 2001 **Proceedings of the tenth international conference on Information and**

knowledge management

Publisher: ACM Press

Full text available:  [pdf\(570.54 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

New applications from the areas of analytical data processing and data integration require powerful features to condense and reconcile available data. As outlined in [1], the general concept of grouping and aggregation appears to be a fitting paradigm for a number of these issues, but in its common form of equality based groups or with current extensions like simple user-defined functions to derive group-by values on a per tuple basis and restricted aggregate functions a number of problems remain ...

4 [Rewriting queries with arbitrary aggregation functions using views](#)



Sara Cohen, Werner Nutt, Yehoshua Sagiv

June 2006 **ACM Transactions on Database Systems (TODS)**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(294.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The problem of rewriting aggregate queries using views is studied for conjunctive queries with arbitrary aggregation functions and built-in predicates. Two types of queries over views are introduced for rewriting aggregate queries: *pure candidates* and *aggregate candidates*. Pure candidates can be used to rewrite arbitrary aggregate queries. Aggregate candidates can be used to rewrite queries containing aggregate functions definable in terms of a commutative-semigroup operation. For ...

Keywords: View usability, query equivalence, query rewriting

5 [Aggregate predicate support in DBMS](#)

Apostol (Paul) Natsev, Gene Y. C. Fuh, Weidong Chen, Chi-Huang Chiu, Jeffrey S. Vitter
January 2002 **Australian Computer Science Communications , Proceedings of the thirteenth Australasian database conference - Volume 5 ADC '02**, Volume 24 Issue 2

Publisher: Australian Computer Society, Inc., IEEE Computer Society Press

Full text available:  [pdf\(1.57 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we consider aggregate predicates and their support in database systems. Aggregate predicates are the predicate equivalent to aggregate functions in that they can be used to search for tuples that satisfy some aggregate property over a set of tuples (as opposed to simply computing an aggregate property over a set of tuples). The importance of aggregate predicates is exemplified by many modern applications that require ranked search, or top-k queries. Such queries are the norm ...

Keywords: aggregate predicates, nearest neighbor, query optimization


6 [Answering complex SQL queries using automatic summary tables](#)



Markos Zaharioudakis, Roberta Cochrane, George Lapis, Hamid Pirahesh, Monica Urata

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(185.85 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We investigate the problem of using materialized views to answer SQL queries. We focus on modern decision-support queries, which involve joins, arithmetic operations and other (possibly user-defined) functions, aggregation (often along multiple dimensions), and nested subqueries. Given the complexity of such queries, the vast amounts of data upon which they operate, and the requirement for interactive response times, the use of materialized views (MVs) of similar complexity is often mandatory ...

7 The PanQ tool and EMF SQL for complex data management



Damianos Chatziantoniou

August 1999 **Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining**

Publisher: ACM Press

Full text available: [pdf\(517.91 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

8 Research papers: estimation and approximation: Relational confidence bounds are easy with the bootstrap



Abhijit Pol, Christopher Jermaine

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(337.51 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Statistical estimation and approximate query processing have become increasingly prevalent applications for database systems. However, approximation is usually of little use without some sort of guarantee on estimation accuracy, or "confidence bound." Analytically deriving probabilistic guarantees for database queries over sampled data is a daunting task, not suitable for the faint of heart, and certainly beyond the expertise of the typical database system end-user. This paper considers the prob ...

9 Classification: SQL database primitives for decision tree classifiers



Kai-Uwe Sattler, Oliver Dunemann

October 2001 **Proceedings of the tenth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: [pdf\(1.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Scalable data mining in large databases is one of today's challenges to database technologies. Thus, substantial effort is dedicated to a tight coupling of database and data mining systems leading to database primitives supporting data mining tasks. In order to support a wide range of tasks and to be of general usage these primitives should be rather building blocks than implementations of specific algorithms. In this paper, we describe primitives for building and applying decision tree classifi ...

Keywords: SQL-aware data mining, data mining primitives, query operators

10 Querying web metadata: Native score management and text support in databases



Gültekin Özsoy[□]lu, Ismail Sengör Altıngövdé, Abdullah Al-Hamdani, Selma Ayşe Özel, Özgür Ulusoy, Zehra Meral özsoy[□]lu

December 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 4

Publisher: ACM Press

Full text available: [pdf\(737.76 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this article, we discuss the issues involved in adding a native score management system to object-relational databases, to be used in querying Web metadata (that describes the semantic content of Web resources). The Web metadata model is based on topics (representing entities), relationships among topics (called *metalinks*), and importance scores (sideway values) of topics and metalinks. We extend database relations with scoring functions and importance scores. We add to SQL score-manag ...

Keywords: Score management for Web applications

11 Sensor networks: Contour map matching for event detection in sensor networks



Wenwei Xue, Qiong Luo, Lei Chen, Yunhao Liu

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(389.31 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Many sensor network applications, such as object tracking and disaster monitoring, require effective techniques for event detection. In this paper, we propose a novel event detection mechanism based on matching the contour maps of in-network sensory data distribution. Our key observation is that events in sensor networks can be abstracted into spatio-temporal patterns of sensory data and that pattern matching can be done efficiently through contour map matching. Therefore, we propose simple SQL ...

Keywords: contour maps, event detection, patterns, sensor networks

12 Industrial sessions: database internals - II: Hosting the .NET Runtime in Microsoft SQL server



Alazel Acheson, Mason Bendixen, José A. Blakeley, Peter Carlin, Ebru Ersan, Jun Fang, Xiaowei Jiang, Christian Kleiner, Balaji Rathakrishnan, Gideon Schaller, Beysim Sezgin, Ramachandran Venkatesh, Honggang Zhang

June 2004 **Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: pdf(249.27 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

The integration of the .NET Common Language Runtime (CLR) inside the SQL Server DBMS enables database programmers to write business logic in the form of functions, stored procedures, triggers, data types, and aggregates using modern programming languages such as C#, Visual Basic, C++, COBOL, and J++. This paper presents three main aspects of this work. First, it describes the architecture of the integration of the CLR inside the SQL Server database process to provide a safe, scalable, secure, an ...

13 NonStop SQL/MX primitives for knowledge discovery



John Clear, Debbie Dunn, Brad Harvey, Michael Heytens, Peter Lohman, Abhay Mehta, Mark Melton, Lars Rohrborg, Ashok Savasere, Robert Wehrmeister, Melody Xu

August 1999 **Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining**

Publisher: ACM Press

Full text available: pdf(526.62 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

14 Data mining (DM): Embedded predictive modeling in a parallel relational database



A. Dorneich, R. Natarajan, E. Pednault, F. Tipu

April 2006 **Proceedings of the 2006 ACM symposium on Applied computing SAC '06**

Publisher: ACM Press

Full text available: pdf(224.69 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A methodology for embedding predictive modeling algorithms in a commercial parallel database is described; specifically, the parallel editions of IBM DB2 Universal Database, although many aspects of the overall approach can be used with other commercial parallel databases. This parallelization approach was implemented in the Version 8.2 release of DB2 Intelligent Miner Modeling to support a new predictive modeling algorithm called Transform Regression. This database-embedded mining algorithm pro ...

Keywords: data mining, embedded analytics, parallel databases, predictive modeling, transform regression

15 Optimization and evaluation of database queries including embedded interpolation



procedures

Leonore Neugebauer

April 1991 **ACM SIGMOD Record , Proceedings of the 1991 ACM SIGMOD international conference on Management of data SIGMOD '91**, Volume 20 Issue 2

Publisher: ACM Press

Full text available: pdf(1.00 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

16 Online aggregation



Joseph M. Hellerstein, Peter J. Haas, Helen J. Wang

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: pdf(1.92 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Aggregation in traditional database systems is performed in batch mode: a query is submitted, the system processes a large volume of data over a long period of time, and, eventually, the final answer is returned. This archaic approach is frustrating to users and has been abandoned in most other areas of computing. In this paper we propose a new online aggregation interface that permits users to both observe the progress of their aggregation queries and control execution on ...

17 Integrating association rule mining with relational database systems: alternatives and implications



Sunita Sarawagi, Shiby Thomas, Rakesh Agrawal

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available: pdf(2.03 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data mining on large data warehouses is becoming increasingly important. In support of this trend, we consider a spectrum of architectural alternatives for coupling mining with database systems. These alternatives include: loose-coupling through a SQL cursor interface; encapsulation of a mining algorithm in a stored procedure; caching the data to a file system on-the-fly and mining; tight-coupling using primarily user-defined functions; and SQL implementations for processing in the DBMS. We ...

18 Query processing and optimization: Object-relational management of complex



geographical objects

Hans-Peter Kriegel, Peter Kunath, Martin Pfeifle, Matthias Renz

November 2004 **Proceedings of the 12th annual ACM international workshop on Geographic information systems**

Publisher: ACM Press

Full text available: pdf(191.98 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern database applications including computer-aided design, multimedia information systems, medical imaging, molecular biology, or geographical information systems impose new requirements on the effective and efficient management of spatial data. Particular problems arise from the need of high resolutions for large spatial objects and from the design goal to use general purpose database management systems in order to guarantee industrial-strength. In the past two decades, various stand-alon ...

Keywords: data management, object decomposition, object-relational database, spatial

19 Extending complex ad-hoc OLAP



Theodore Johnson, Damianos Chatziantoniou

November 1999 **Proceedings of the eighth international conference on Information and knowledge management**

Publisher: ACM Press

Full text available: pdf(1.36 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Large scale data analysis and mining activities require sophisticated information extraction queries. Many queries require complex aggregation, and many of these aggregates are non-distributive. Conventional solutions to this problem involve defining User Defined Aggregate Functions (UDAFs). However, the use of UDAFs entails several problems. Defining a new UDAF can be a significant burden for the user, and optimizing queries involving UDAFs is difficult because of the "black box" nature of the aggregation function.

20 [Ranking: Ranking objects based on relationships](#)



Kaushik Chakrabarti, Venkatesh Ganti, Jiawei Han, Dong Xin

June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: pdf(278.59 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In many document collections, documents are related to objects such as document authors, products described in the document, or persons referred to in the document. In many applications, the goal is to find these objects that best match a set of keywords. However, the keywords may not necessarily occur in the target objects; they occur only in the documents. For example, in a product review database, a user might search for names of products (say, laptops) using keywords like "lightweight" and " ...

Keywords: aggregation, early termination, keyword search, named entities, ranking, relationships, top-k queries

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:



[Adobe Acrobat](#)



[QuickTime](#)



[Windows Media Player](#)



[Real Player](#)

- ☒ Search only in Engineering, Computer Science, and Mathematics.
☐ Search in all subject areas.

Scholar All articles Recent articles

Results 1 - 10 of about 19,800 for aggregate class interface. (0.17 seconds)

All Results

P Szekely

D Clark

R Campbell

M Buchi

M Mattsson

A user interface toolkit based on graphical objects and constraints

P Szekely, B Myers - ACM SIGPLAN Notices, 1988 - portal.acm.org

... figure 4. The figure shows two check-boxes such as those used in the Macintosh user

interface, and a Coral ... (Herr-CO Check-Box (: super-classes **Aggregate-Yith-go** ...

Cited by 48 - Related Articles - Web Search

Framework composition: problems, causes and solutions - group of 11 »

M Mattsson, J Bosch - Technology of Object-Oriented Languages and Systems, 1997. ... - ieeeexplore.ieee.org

... no name clashes occur, this solution is very suitable, but it results in some imple-
mentation overhead as for each operation in the **class interface** that has ...

Cited by 38 - Related Articles - Web Search

A++/P++ array classes for architecture independent finite difference computations - group of 2

»

R Parsons, D Quinlan - 2. Conference on object-oriented numerics (OONSKI 94, 1994 - osti.gov

... to implement parallel support using the A++ array **class interface**. Page 6. ... expression,

the time for the array **class** library without **aggregate** operations ...

Cited by 33 - Related Articles - View as HTML - Web Search

Supporting real-time applications in an Integrated Services Packet Network: architecture and

... - group of 33 »

DD Clark, S Shenker, L Zhang - Applications, Technologies, Architectures, and Protocols for ..., 1992 -
portal.acm.org

... attempted in our proposal to produce an **interface** which is ... 2.1 A **Class** of Real-Time

Applications In the discussion that follows, we focus on a particular **class** ...

Cited by 668 - Related Articles - Web Search - BL Direct

Choices (**class** hierarchical open **interface** for custom embedded systems) - group of 2 »

R Campbell, G Johnston, V Russo - ACM SIGOPS Operating Systems Review, 1987 - portal.acm.org

... The **aggregate** of the Spaces in 12 ... schemes can be built by extending the **class** hierarchy ...

An **interface** compiler for C++ enriches the possible communication schemes ...

Cited by 75 - Related Articles - Web Search

Generic wrappers - group of 11 »

M Buchi, W Weck - Proceedings of ECOOP, 2000 - Springer

... A Java **interface** corresponds to a fully abstract **class** in ... t know whether an instance

of her **class** will wrap a ... of the static wrappee type (2). An **aggregate** of a ...

Cited by 57 - Related Articles - Web Search - BL Direct

[book] The unified modeling language - group of 8 »

J Rumbaugh, I Jacobson - 1996 - microgold.com

... Only one side (at most) of an association can be an **aggregate**. ... **Interface** Designates

the **interface** of a **class** or a package, consisting of a ...

Cited by 2587 - Related Articles - View as HTML - Web Search - Library Search

A Language Facility for Designing Database-Intensive Applications - group of 6 »

J MYLOPOULOS, PA BERNSTEIN, HKT WONG - ACM Transactions on Database Systems, 1980 -

portal.acm.org

... The DBMS provides an **interface** into which the database operations of the IIS can

Cited by 272 - Related Articles - Web Search - Library Search

application - group of 4 »

S Parkes, JA Chandy, P Banerjee - Proceedings of the 1994 ACM/IEEE conference on ..., 1994 - portal.acm.org

open, any of the system **aggregate** types may be used as a base **class** in an ...

Cited by 22 - Related Articles - Web Search

ObjectStore, and ... - group of 2 »

V Soloviev - ACM SIGMOD Record, 1992 - portal.acm.org

and others. The ONTOS SQL **interface** provides SQL access to ONTOS ...

[Cited by 25](#) - [Related Articles](#) - [Web Search](#)

Goooooooooooooogle ►

Result Page: 1 2 3 4 5 6 7 8 9 10 **Next**

aggregate class interface

[Google Home](#) - [About Google](#) - [About Google Scholar](#)

©2006 Google